

ABSTRACT OF THE DISCLOSURE

A method for detecting and repairing cloud splits in a distributed system such as a peer-to-peer (P2P) system is presented. Nodes in a cloud maintain a multilevel cache of entries for a subset of nodes in the cloud. The multilevel cache is built on a circular number space, where each node in the cloud is assigned a unique identifier (ID). Nodes are recorded in levels of the cache according to the distance from the host node. The size of the cloud is estimated using the cache, and cloud-split tests are performed with a frequency inversely proportional to the size of the cloud. Cloud splits are initially detected by polling a seed server in the cloud for a node N having an ID equal to the host ID + 1. The request is redirected to another node in the cloud, and a best match for N is resolved. If the best-match is closer to the host than any node in the host's cache, a cloud split is presumed. The cloud split is repaired by flooding the host's address to the newly found node and sending repair messages to nodes in the host's top cache level. Each node receiving a repair message repeats a similar process, and sends repair messages to nodes in its next lower cache level.